

42390.P9919

Patent

UNITED STATES PATENT APPLICATION

FOR

**Apparatus and Method For
Facilitating Access To Network Resources**

INVENTORS:

Kingsum Chow
Colin Cunningham
Thomas Holman

00327-5224250

Prepared by

Steven D. Yates
Reg. No. 42,242
(503) 264-6589

Express Mail mailing label number: EL034436510US

Apparatus and Method For Facilitating Access To Network Resources

Field of the Invention

The invention generally relates to providing network resources to clients, and more particularly to providing mirrored network resources with multiple network hosts, where a requesting client is automatically directed to a network host having a more efficient communication channel with the client.

Background

With the widespread availability of intranets and the Internet in the home and workplace, network traffic has become increasingly congested, leading to increasing client delays in obtaining desired network resources.

In an effort to avoid such delays, a common technique is to host a network site, e.g., a web site or other network resource, on multiple network hosts in different geographic areas. Thus, a network site may be hosted in different countries and localities within the countries. An incoming client networking connection is then manually or automatically redirected to a host geographically closest to the client.

In a manual environment, on contacting one host, the host returns to the client a network resource, such as a web page, providing alternative hosts for the client. For example, assuming a client receives a web page, the page may contain hyperlinks to the available network hosts for the contacted site.

In an automatic redirection environment, when the network site's network name, e.g., a Uniform Resource Locator (URL), is resolved by a Domain Name Server (DNS), rather than having the DNS return an established (or static) mapping of a network address for the network site's name, instead the DNS returns the network address of whichever network host is geographically closest to the client. It is assumed that host addressing conforms to naming devices on a network, with DNS supplying an address, such as a Transmission Control Protocol/Internet Protocol (TCP/IP) address, for the named device. The theory is that geographically distant network hosts will have longer network delays, while close network hosts will have short delays. Unfortunately, such an arrangement does not always result in the client being connected to the network host having the most efficient connection to the client, e.g., the network host providing the fastest data response times to the client.

Brief Description Of The Drawings

The features and advantages of the present invention will become apparent from the following detailed description of the present invention in which:

FIG. 1 illustrates a generalized network environment in accordance with one embodiment of the invention.

FIG. 2 is a flowchart illustrating one embodiment for a client to obtain a network site's resources by determining and directing the client connection to an efficient source for the desired network resources.

FIG. 3 illustrates one embodiment for redirecting the client to the more efficient communication source.

FIG. 4 illustrates another embodiment for redirecting the client to multiple more efficient communication sources.

FIG. 5 is a flowchart illustrating one embodiment for determining an efficient communication source to which to direct the client.

FIG. 6 illustrates an exemplary table for tracking clients and their communication efficiencies for different network hosts.

FIG. 7 illustrates a generalized approach to determining efficiency ratings.

FIG. 8 illustrates a suitable computing environment in which certain aspects of the invention may be implemented.

Detailed Description

FIG. 1 illustrates a generalized network environment in accordance with one embodiment of the invention.

As illustrated, a United States based Client 1 **100** and a Europe based Client 2 **102** are in communication with a network **104**, such as an intranet, the Internet, or other networking environment (e.g., wireless, satellite, etc.). Although a single network is illustrated, it will be appreciated that there may be many individual networks, wired and wireless, that are interconnected to form the illustrated network. Clients **100**, **102** seek network resources, such as web pages or other data, from a network site **106**. To facilitate access to the network site's data, some or all of the site's network resources are "mirrored" or otherwise made available through other network hosts **108**, **110**, **112**.

As illustrated, the other mirroring network hosts may include, for example, a United States based Network Host **108**, a Europe based Network Host **110**, and an Asia

based Network Host **112**. As indicated by the ellipses, there may be many more hosts in other localities, each mirroring some or all of Original Network Site resources. As used herein and the claims that follow, unless indicated otherwise either explicitly or implicitly through context, the phrase “network host” may collectively reference the Original Network Site **106** as well as the mirroring network hosts **108, 110, 112**.

When Client 1 **100** seeks to access the network site **106**, it is assumed that Client 1 provides the name of the Original Network Site to a name resolution service, such as Domain Name Server (DNS) 1 **114**. In response, DNS 1 returns a network address for the Original Network Site **106**, or that of the mirroring network hosts **108, 110, 112**, depending on the resolution strategy in use.

Assuming a prior art geographic-base resolution strategy, if the United States based Network Host **108** is geographically closest to Client 1 **100**, then when Client 1 **100** attempts to resolve the Original Network Site’s **106** network name, DNS 1 **114** will direct Client 1 to the United States Network Host. Client 1 then communicates directly with the United States Network Host to obtain network resources. Similarly, if Client 2 **102** tries contacting the Original Network Site, since Client 2 is in France, with geographic-based resolution, DNS 2 **116** returns the network address for the Europe based Network Host **110**.

One problem with the geographic based approach, however, is that it assumes geographic proximity equates with communication efficiency. Unfortunately this may not be the case. For example, local and regional network congestion, problems in network hardware, incorrectly configured equipment, etc., can result in a nearest of the network hosts being a less efficient data source than one of the other network hosts.

Or, a particular remote network host might be reachable over a higher-speed data pathway, such as an Internet "backbone" (e.g., a fast Synchronous Optical Network (SONET) Optical Carrier (OC)). Thus, the result provided by a geographical-based resolution DNS may be sub-optimal.

FIG. 2 is a flowchart illustrating one embodiment for Client 1 **100** (FIG. 1) to select between available network hosts **106-112** for obtaining the Original Network Site's **106** network resources according to the present invention. In one embodiment, in order to ensure that Client 1 is connected to a network host providing a more efficient communication session, response times for receiving network data from the various network hosts are measured, and the host having fastest response time is utilized.

A first operation is for Client 1 **100** to submit **200** the Original Network Site's name to its resolution service, e.g., DNS 1 **114**. DNS 1 then returns **202** a network address corresponding to the Original Network Site. In one embodiment, DNS 1 returns an address for a geographically closest network host as in the prior art. In an alternate embodiment, DNS 1 simply returns the network address for the Original Network Site. Once Client 1 has a resolved network address, it then connects **204** to the provided network address. In response, Client 1 then receives **206** initial network resources from the network host.

In one embodiment, before returning the network resource to Client 1 **100**, the contacted network host determines the most efficient communication pathway to Client 1, e.g., it determines the network host with which Client 1 ought to be communicating. Towards this end, the Original Network Site **106** and mirroring network hosts **108, 110**,

112 are inspected (FIG. 5) to determine an efficiency rating tracked by each host regarding communication with Client 1. The network host (or hosts) identified as being most efficient is then encoded within the network resource returned to Client 1.

Based on this encoding, Client 1 then continues its retrieval 208 of network resources according to the embedded reference. It will be appreciated that the resource can be of any data type, including web page code, audio data, video data, a database, operational commands / directives for controlling Client 1, another data type allowing for embedded links / references to other resource locations, or some combination of these data types. Thus, for example, a network resource retrieved from a first host may be a Joint Photographic Experts Group (JPEG) graphics file containing embedded tags linking to a more efficient second host for processing the image.

FIG. 3 illustrates one embodiment in which the returned network resource is a web page 300, and the encoded reference is an embedded link 302 to a resource located on the network host 106-112 that was identified as being most efficient.

In this embodiment, Client 1 100 contacts the network site indicated by DNS 114 and receives a web page in which its web page links, rather than referencing the originator of the web page, instead directs Client 1 towards the network host identified as being most efficient.

In another embodiment (not illustrated), the returned network resource is a redirection request to redirect Client 1 100 to communicate directly with the efficient communication source. Common examples of redirection commands include HyperText Markup Language (HTML) redirection commands, Java/JavaScript code to

redirect a browser, and Common Gateway Interface (CGI) / Perl scripts to redirect Client 1. In this embodiment, rather than Client 1 receiving a substantive network resource from the contacted network host, such as an entire web page, instead Client 1 **100** receives minimal data required to effect a redirection of its communication.

FIG. 4 illustrates another embodiment, in which the returned network resource **400** from the network host identified by a DNS includes multiple encoded references to multiple network hosts, with each encoded reference providing a link to the most efficient communication source for the type of resource being linked.

For example, if the returned resource is a web-page, and the web page is to include links **402** to streaming audio and/or visual data, links **404** to electronic commerce (e-commerce), links **406** to database resources, links **408** to telephony operations, etc., then these embedded web page links **402-408** can direct Client 1 to the network hosts **106-112** having highest efficiency for that resource.

In the case of streaming media or other such data types, the link may automatically initialize / start delivery of the streaming media. Thus, with respect to Client 1 **100**, U.S. Network Host **108** (FIG. 1) may be, for example, the most efficient source for audio data (links **402**), while the Original Network Site **106** may be the most efficient source for engaging in e-commerce transactions (links **404**).

FIG. 5 is a flowchart illustrating one embodiment for determining which of the Original Network Site **106** and mirroring Network Hosts **108, 110, 112** is the most efficient source for retrieving desired network resources. In one embodiment, each

possible network host (e.g., hosts **106-112**) for handling a client's (e.g., Client 1 **100**) request for network resources maintains a table (see FIG. 6), database, or other data structure tracking known efficiency ratings for communication with the client.

Assume Client 1 **100** attempts to locate **500** the Original Network Site **106** to retrieve certain network resources, such as a web page. As discussed above, assuming geographic resolution, Client 1's DNS returns **502** the network address of the United States Network Host **108** due to it being geographically closest to Client 1.

Client 1 then establishes **504** a network connection with the United States Network Host **108**. In response to this contact, the United States Network Host looks up **506** Client 1 in its tracked efficiency ratings, and determines **508** whether another network host **106, 110, 112** is a more efficient source for Client 1's resource request. Assuming the Europe Network Host **110** is most efficient, then the U.S. Network Host constructs **510** a return network resource, in this case a web page, containing encoded references to the Europe Network Host.

FIG. 6 illustrates an exemplary table **600** for tracking clients and their communication efficiencies for different network hosts. It is assumed that a copy of this table is available to all network hosts, or that the table is shared among all hosts.

In the illustrated embodiment, tracked communication efficiency data is indexed according to a client's network address, e.g., it's Internet Protocol (IP) address if a TCP/IP network is used. As illustrated, addresses are stored in the first column **602** of the table. The remaining columns **604-610** store tracked efficiency values indicating efficiency of communication between a particular client and the network hosts.

In one embodiment, each table entry in columns **602-610** include sub-entries storing reliability ratings (not shown) that can be used to weight or rank stored efficiency ratings. Thus, a client ay be directed to a host having a lesser efficiency, but higher reliability, than another host. It will be appreciated that the illustrated table **600** comprises an efficient and compact data structure, since it provides a 1:1 mapping between table columns and the number of network hosts **106-112** (FIG. 1) (e.g., data centers), and the number of rows may be limited (if desired) according to thresholds or caching techniques.

FIG. 7 illustrates a generalized approach to determining efficiency ratings. Generally, the efficiency ratings stored in the FIG. 6 table can be an average, moving average, or other statistic or heuristic measurement of actual communication performance with the client, and therefore account for real-time disturbances in communication data paths between the client and network hosts.

In one embodiment, the communication efficiency ratings stored in the tracking data correspond to measured communication delays between the client and a network host **106-112** maintaining the tracking data. In one embodiment, efficiency ratings factor in past and/or predicted reliability of a host.

As illustrated, a client contacts **700** a first network host. In response the client receives **702** a network resource from the first network host, where the received network resource contains embedded references which cause the client to request **704** additional resources from the first network host so as to allow the first network host to determine **706** communication efficiency with the contacting client. In one embodiment,

the network resource also directs the client to request additional resources from other network hosts so that they can also determine communication efficiency.

For example, assuming the resource initially received from the first network host is a web page, the web page has embedded links to additional resources, such as to graphics images. Knowing the sizes of the additional resources, the first network host can measure the time delay between the requests for the additional resources to predict the communication efficiency between the client and the first network host. The predicted value is then stored **708** in a FIG. 6 table. As noted above, the web page may be configured to cause the client to perform similar resource requests from other network hosts so they can also determine communication efficiency values.

In one embodiment, the tracked efficiency rating incorporate status information provided by network and server monitoring systems / services, so that systems determined to have problems will be updated as having very inefficient values (or a special value, such as a negative number, indicating complete unavailability). Thus, if a particular server is down, this failure impacts its communication efficiency with the client, which in turn causes a different network host to be determined as being the most efficient source for the client's desired resources.

In addition, determining an efficient host based on real time interactions between a client and each network host providing desired network resources allows one to also perform load balancing of client contact requests. A host having a high load will automatically produce a less efficient response time, causing the client's communication efforts to be directed elsewhere.

Tracked efficiency ratings may be revised on a periodic basis. In one embodiment, the tracked efficiency ratings are revised as a function of the number of data transactions occurring with Client 1. For example, assuming Client 1's desired network resource is web page data, then a small percentage of the web transactions, e.g., 0.01%, can be configured to cause new efficiency ratings to be measured. Measured and revised efficiency ratings are propagated to all network hosts **106-112**.

FIG. 8 and the following discussion are intended to provide a brief, general description of a suitable computing environment in which certain aspects of the illustrated invention may be implemented.

An exemplary system for implementing the invention includes a computing device **800** having system bus **802** for coupling various computing device components. Typically, attached to the bus are non-programmable and programmable processors **804**, a memory **806** (e.g., RAM, ROM), storage devices **808**, a video interface **810**, and input/output interface ports **812**. Storage devices include hard-drives, floppy-disks, optical storage, magnetic cassettes, tapes, flash memory cards, memory sticks, digital video disks, and the like.

The invention may be described by reference to different high-level program modules and/or low-level hardware contexts. Those skilled in the art will realize that program modules can be interchanged with low-level hardware instructions. Program modules include procedures, functions, programs, components, data structures, and the like, for performing particular tasks or implementing particular abstract data types. Modules may be incorporated into single and multi-processor computing devices,

Personal Digital Assistants (PDAs), cellular telephones, and the like. Thus, the storage systems and associated media can store data and executable instructions for the computing device.

The computing device is expected to operate in a networked environment using logical connections to one or more remote computing devices **814, 816** through a network interface **818**, modem **820**, or other communication pathway. Computing devices may be interconnected by way of a network **822** such as an intranet, the Internet, or other network. Modules may be implemented within a single computing device, or processed in a distributed network environment, and stored in both local and remote memory. Thus, for example, with respect to the illustrated embodiments, assuming computing device **800** is Client 1 **100** (FIG. 1) seeking to obtain a network resource from an Original Network Site **106**, then remote devices **814, 816** may respectively be the Original Network Site and the United States Network Host **108** mirroring some or all of the Original Network Site's network resources.

It will be appreciated that remote computing devices **814, 816** may be configured like computing device **800**, and therefore include many or all of the elements discussed for computing device. It should also be appreciated that computing devices **800, 814, 816** may be embodied within a single device, or separate communicatively-coupled components, and may include or be embodied within routers, bridges, peer devices, web servers, and application programs utilizing network application protocols such as the HyperText Transfer Protocol (HTTP), File Transfer Protocol (FTP), and the like.

Having described and illustrated the principles of the invention with reference to illustrated embodiments, it will be recognized that the illustrated embodiments can be modified in arrangement and detail without departing from such principles.

And, even though the foregoing discussion has focused on particular embodiments, it is understood that other configurations are contemplated. In particular, even though expressions such as "in one embodiment," "in another embodiment," or the like are used herein, these phrases are meant to generally reference embodiment possibilities, and are not intended to limit the invention to particular embodiment configurations. As used herein, these terms may reference the same or different embodiments, and unless implicitly or expressly indicated otherwise, embodiments are combinable into other embodiments. Consequently, in view of the wide variety of permutations to the above-described embodiments, the detailed description is intended to be illustrative only, and should not be taken as limiting the scope of the invention.

What is claimed as the invention, therefore, is all such modifications as may come within the scope and spirit of the following claims and equivalents thereto.